Towards a Definition of Robustness for Market-Style Open Multi-Agent Systems

Michael Schillo[1], Hans-Jürgen Bürckert[2], Klaus Fischer[2], Matthias Klusch[2]

schillo@virtosphere.de, {hjb, kuf, klusch}@dfki.de

Multi-agent Systems Group[1]
Universität des Saarlandes
Im Stadtwald, Geb. 43.8
D- 66123 Saarbrücken

Multi-agent Systems Group[2]
DFKI GmbH
Im Stadtwald, Geb. 43.8
D- 66123 Saarbrücken

Keywords

Multi-agent simulation, social agents, social reasoning, organisation of agent societies.

1. MOTIVATION

Multi-agent systems (MAS) are an integral part of current (D)AI research. Their success and the success of agent technology is in the majority of the literature ascribed to three major advantages, one of which is that multi-agent systems are robust systems (cf. [1, 9]). Although the literature treats robustness like an inherent feature of MAS (which it is not), there is hardly any discussion on what robustness actually means (a rare exception is the work of Kaminka and Tambe [7] and Klein and Dellarocas [3]). Just like any other artificial complex system, multi-agent systems need to be specifically designed to be robust. Compared to conventional computer science, the issue of robustness in MAS is different. Most computer science systems are transformational systems, which means they compute a function on some input. Here, techniques for ensuring robustness exist (e.g. cf. [6]). However, the most interesting multi-agent systems are open systems, which do not explicitly compute a function (e.g. looking at the most predominant example of an open system, namely the Internet, as computing a function certainly misses the point).

2. ROBUSTNESS IN MAS

Robustness in multi-agent systems is more than introducing redundancy. Redundancy will not solve problems such as malicious agents in an open system, communication and information overload. A broader conception of robustness is required. From our point of view, *robustness* can only be defined in relation to some definition of *performance measure*. Robustness is the ability of a system to maintain "safety-responsibilities" [10], even though events happen that are able to disturb the system. So these safety-responsibilities must be defined, as well as it must be defined when the system performs well and how this can be quantified. When looking at an electronic market, we can for example identify the following three performance criteria. Firstly, how fast can a customer find someone who provides a desired service (speed of match-making) and secondly, do both parties, customers and producers, meet their needs, i.e. can producers earn enough to maintain their service and can customers find the service they need. Thirdly, the quality of service provided is of importance (product according to specification, on-time delivery of service, drop-outs, etc.). We define robustness *quantitatively* by the expected drop of the performance measure in four perturbation scenarios (i) **increase of population size**, (ii) **change of task profile** over time, (iii) malicious agent **intrusion**, and (iv) **drop-outs** of agents. Robustness can then be defined and measured by how much decrease of the performance measure will be the result of doubling the population size, five percent random drop-outs, etc. If the limit for calling a system robust is defined to be five percent, then there is a clear-cut *qualitative* definition of whether a system is robust. Mastering the perturbation scenarios corresponds to providing the following four properties:

Scalability: How does the system react in terms of performance if the size of the agent society is increased by a certain percentage. Specifically, this requires that patterns of interaction can react to the increased size of the community. Possible strategies are choosing different protocols, employing matchmakers, or organisation of participating parties to bigger entities (cf. the literature on *holons*, e.g. [5]).

Flexibility: If any change in the environment occurs and the safety responsibilities cannot be maintained at the moment, is the system able to recover? How fast can the system recover to such disturbances and how fast can agents adopt their models of others (some kind of modelling is present in almost any multi-agent system, if not it is interesting to investigate how it deals with lack of knowledge about others *in combination* with scalability).

Resistance: It is interesting the effect of malicious behaviour of agents in the community of agents (namely lying in communication about facts in order to manipulate knowledge of agents, abusing protocols). Work on trust and betrayal in agent societies can be found in [8], which deals with the issue of betrayal in communication about others. In this work we showed that it is possible to increase the robustness of systems towards malicious agents by using trust as a complex mechanism for evaluating agents and excluding malicious agents from interaction in the population.

Drop-out safety: The reasons for agents halting execution are manifold, handling these situations is difficult. Approaches in this context are "shadow agents" i.e. agents that monitor other agents and replace them [11] or adapt their plans accordingly [7] and market-based approaches where requests are announced in contract-net protocol fashion and drop-outs are compensated by other agents getting their bids accepted. This requires a theory of delegation (see [4] and below).

3. HOW TO OBTAIN ROBUSTNESS

In the following we want to extend a theory of delegation that will help provide the above four properties in multi-agent systems. This extension is inspired by the Habitus-Field theory of sociologist Pierre Bourdieu [2], which is of importance to DAI in that it helps to extend individual (psychological) theories to distributed (social) theories. From Bourdieu's theory, we can derive four mechanisms of interaction of natural distributed systems that explain the stabilisation as well as the flexibility of a group or team. The interplay of theses mechanisms contributes enormously to the robustness of social systems.

3.1 Types of Operation

Delegating tasks to other agents is not new to MAS research, research on task-oriented domains has for a long time been involved in how to distribute the right task to the right agent. But this model of delegation is restricted to two kinds of settings: settings where agents are benevolent, i.e. they are all designed to share common goals, and settings where agents simulate authority relationships. Neither of these settings applies in an electronic market. Here, a further mode is present, namely the negotiation of delegation. This means that agents decide on a case by case basis whether they delegate a task and to whom. Recent work on delegation (see [4] for an extensive treatment), has shown that delegation is a complex concept and at least in semi-open systems to be very relevant to multi-agent systems research. This point is stressed also by Bourdieu for human societies, when attributing a central role to delegation. The mechanism of delegation makes it possible to pass on tasks to other individuals and furthermore, allows specialisation of these individuals for certain tasks (functional differentiation and role performance). Thus, we need to differentiate *two types of operation: task delegation*, which is the delegation of a sequence of (autistic) goals to be achieved and social *delegation*, which does not consist of creating a solution or a product but in representing, for example a group or organisation in a negotiation. Both types of operation are essential for organisations, as their strength is that they do not rely on a particular individual, but on somebody able to perform this role.

3.2 Mechanisms for Delegation

It is important to note that finding the right individual for delegating a task to is no trivial task. We observe four different mechanisms to determine a delegate: (i) a very well known mechanism is *voting*, whereby a group of equals determines by some voting mechanism (majority, two thirds, etc.) one of them to be the delegate. (ii) Authority is another well-known mechanisms, it represents the method of organisation used in *distributed problem solving*. (iii) *economic exchange* is a standard mode in markets. A good is exchanged for money, while the involved parties assume that the value of both is of appropriate similarity, i.e. the delegate is being paid for doing the delegated task or representation (iv) *social exchange:* According to Bourdieu, this is not the only kind of exchange in social settings. Rather, gifts are being given (think of a gift in an abstract way: this can for example mean that someone accepts dissimilarity in an economic exchange) and favours are being done, all in expectation of either reciprocation or refusal of reciprocation. Both are indications to the involved parties about the state of their relationship. This kind of exchange entails risk, trust, and the possibility of conflicts (continually no reciprocation) and the need for an explicit management of relationships. The aim of this is to accumulate "social capital" that may pay-off in the future, much like a storage that can be used up in times of scarcity. Obviously, this kind of "storage" is also a source of robustness. Note that these four mechanisms work for both types of operation: economic exchange can be used for social delegation and voting for task delegation as well as vice versa.

3.3 DELEGATION AND ROBUSTNESS

We believe that the two types of operation combined with the four mechanisms provide the basis to achieve the four properties necessary to master the perturbation scenarios. Social delegation supports scalability of multi-agent systems in that it structures groups of agents and reduces communication. With the aid of task delegation we believe that multi-agent systems can achieve the flexibility to react to changing task profiles. Social exchange and the entailed concepts of trust and risk deal with norm-breaking agents. The thoroughly applied concept of delegation can provide the mechanism necessary to deal with drop-outs. Possibly, the described types of operation and mechanisms are not complete, but we believe they are necessary ingredients for robust multi-agent systems.

4. REFERENCES

[1] Bond, A. H. and Gasser, L. (1988). *Readings in Distributed Artificial Intelligence*, Morgan Kaufmann.

[2] Bourdieu, P., Wacquant, L. (1992). *An Invitation to Reflexive Sociology*, Chicago, Polity Press.

[3] Dellarocas, C., and Klein, M. (2000). An Experimental Evaluation of Domain-Independent Fault Handling Services in Open Multi-Agent Systems. In Proceedings of the 4th International Conference on Multi-Agent Systems, Boston, MA.

[4] Castelfranchi, C. and Falcone, R. (1998). Towards a theory of delegation for agent-based systems. *Robotics and Autonomous Systems*, vol. 24, pp. 141-157.

[5] Gerber, C., Siekmann, J., and Vierke, G. (1999). Flexible Autonomy in Holonic Agent Systems. In *Proc. of AAAI Spring Symposium 1999*.

[6] Ginsberg, M., Parkes, A., and Roy, A. (1998). Supermodels and Robustness. Proceedings of the Fifteenth National Conference on Artificial Intelligence AAAI 98, pp. 334-339.

[7] Kaminka, G. and Tambe, M. (1998). What is wrong with us? Improving robustness through social diagnosis. In *Proceedings of the 15th National Conference on Artificial Intelligence (AAAI-98).*

[8] Schillo, M., Rovatsos, M. and Funk, P. (2000). Using trust for Detecting Deceitful Agents in Artificial Societies. In Castelfranchi, C., Falcone, R., Firozabadi, B., and Tan, Y. (2000). *Trust in Agents*, Special Issue of Applied Artificial Intelligence, vol. 14 (8).

[9] Weiß, G. (1999). *Multi-agent Systems: A Modern Approach to Distributed Artificial Intelligence.* MIT Press.

[10] Wooldridge, M., Jennings, N., und Kinny, D. (1999). A Methodology for Agent-Oriented Analysis and Design. In *Proceedings of the 3rd International Conference on Autonomous Agents, AA99*, pp. 69-76.

[11] Zinnikus, I. and Funk, P. (To appear). Towards Self-stabilizing Agent-based Information Sources. Technical Memo, DFKI, Saarbrücken.